

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : G06F 17/10	A1	(11) International Publication Number: WO 00/14649 (43) International Publication Date: 16 March 2000 (16.03.00)
(21) International Application Number: PCT/GB99/02715 (22) International Filing Date: 20 August 1999 (20.08.99) (30) Priority Data: 98307088.9 3 September 1998 (03.09.98) EP (71) Applicant (for all designated States except US): WALLAC OY [FI/FI]; P.B. 10, FIN-20101 Turku (FI). (72) Inventors; and (75) Inventors/Applicants (for US only): HOBSON, Anthony [GB/GB]; Cavendish Laboratory, University of Cambridge, Madingley Road, Cambridge CB3 0HE (GB). LASENBY, Anthony [GB/GB]; Cavendish Laboratory, University of Cambridge, Madingley Road, Cambridge CB3 0HE (GB). (74) Agent: GILL JENNINGS & EVERY; Broadgate House, 7 Eldon Street, London EC2M 7LM (GB).		(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i>
(54) Title: SIGNAL PROCESSING (57) Abstract A method of reconstructing a signal from a given set of data, with a prediction function representing a predictable effect on the signal, and a noise function representing unpredictable noise. The method comprises the steps of altering the coordinate basis of the data and signal from an original coordinate basis in order to produce a prediction function having a reduced set of variables, performing a Bayesian reconstruction capable of operation of positive, negative, and complex signal values to produce a reconstruction signal, and converting the reconstruction signal back into the original coordinate basis to generate a signal.		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

SIGNAL PROCESSING

The present invention relates to the reconstruction of signals. There are many applications, such as radar, sonar, acoustic data, spectroscopy, geophysical, and image signal processing, in which it is desirable to reconstruct signals from given data.

In many of these situations the effect on the signals of noise and the particular characteristic of the system generating the signal are known or can be approximated using appropriate mathematical models. In these situations, Bayesian reconstruction methods have often been applied to reconstruct the signal from given data. These methods can work well. For example one Bayesian reconstruction approach, known as the maximum entropy method, is known to work well. Usually, the maximum entropy method (MEM), is only be applied to the reconstruction of signals that are strictly non-negative, and which are not complex. Given this, it is generally not possible to change the basis in any vector space representing the data during any reconstruction. Being unable to do this results in computationally very expensive reconstruction processes having to be employed, because of the requirement for very large calculations having an extremely large number of variables to be determined.

For example, employment of such an MEM to a stack of twenty images from a microscope takes in the region of fifty minutes using a standard INTEL™ Pentium™ 200MHz processor.

The present invention is directed towards improving the reconstruction of signals from given data so that such reconstruction can be performed within a reduced time frame.

According to the present invention there is provided a method of reconstructing a signal from a given set of data, with a prediction function representing a predictable

effect on the signal, and a noise function representing unpredictable noise, the method comprising the steps of:

altering the coordinate basis of the data and signal from an original coordinate basis in order to produce a prediction function having a reduced set of variables;

performing a Bayesian reconstruction capable of operation of positive, negative, and complex signal values to produce a reconstruction signal; and

converting the reconstruction signal back into the original coordinate basis to generate a signal.

The Bayesian reconstruction may be performed using a Fourier basis, or may use a wavelet basis.

The Bayesian reconstruction may employ the maximum entropy method, in which case the method may employ an evaluation parameter, α , which may be determined from a prior reconstruction, set at a fixed value, or determined during the reconstruction step.

The signal to be reconstructed may be an image signal, or may be a radar, sonar, or acoustic data signal. Alternatively, it may be a signal from spectroscopy or a geophysical signal.

By employing the method of the present invention, an example stack of 20 microscope images takes approximately 45 seconds to reconstruct using a standard INTELTM PentiumTM 200 MH² processor.

One example of the present invention will now be described.

Bayesian reconstruction methods have been applied to numerous problems in a wide variety of fields. In their standard form, however, they can be very computationally intensive, since they generally require the numerical maximisation of a complicated function of many variables. For example, in image reconstruction problems, it is not unusual for the number of variables to be $\sim 10^6$. Furthermore, one of the most popular Bayesian reconstruction algorithms is the maximum-entropy method (MEM), which can only be applied to the reconstruction of signals that are strictly non-negative (see below). The method can, however, be extended to signals that can take both positive and negative values. We develop the MEM approach so that it can be applied to the reconstruction of signals that can take positive, negative or complex values. As a result, this enables the use of similarity transformations in the reconstruction algorithm so that calculations can be performed in an alternative "basis" this is more appropriate to the problem under consideration. Specifically, the basis is chosen so that signal is reconstructed by performing a large number of numerical maximisations of low dimensionality, rather than a single maximisation of high dimensionality. This results in a significant increase in speed. Indeed, in the example outlined below, the speed of the reconstruction algorithm is increased by a factor of about 100.

The standard Bayesian reconstruction techniques mentioned above are typically applied to a given data set $d(y)$ in order to reconstruct some underlying signal $s(x)$. Here, y denotes the space over which the data are defined and x denotes the space of the signal, which may in general be distinct from x and need not have the same number of dimensions. For example, given some data stream $d(t)$, which consists of the measured values of some quantity as a function of time t , we may wish to reconstruct the two-

dimensional spatial variation of some other quantity (or signal) $s(x,y)$.

For most measurements it is convenient (or necessary) to digitise the data and the signal (for example if either is to be stored/analyzed on a digital computer). We may therefore denote the data by the vector d with N_d components, where N_d is the number of data samples. Similarly, we denote the signal by the vector s of length N_s , where N_s is the number of points at which we wish to reconstruct the signal.

In general, we may express the data vector as some function θ of the signal vector, i.e.

$$d = \theta(s)$$

The function θ can be non-linear and specifies the effect of the measuring apparatus on the signal that we wish to reconstruct. It is customary to divide this function into the predictable effect of the measuring apparatus on the signal and the stochastic noise part due to inherent inaccuracies in the measurement process. In this case, we may write

$$d = \Phi(s) + \epsilon, \quad (1)$$

Where Φ denotes the predictable response of the apparatus to the signal and ϵ is a vector of length N_d that contains any stochastic noise contributions to the data.

The Bayesian approach to reconstructing the signal is to calculate the estimator \hat{s} that maximises the posterior probability $\Pr(s|d)$. This is given by Bayes' theorem as

$$\Pr(s|d) = \frac{\Pr(d|s) \Pr(s)}{\Pr(d)},$$

where $\Pr(d|s)$ is the likelihood of obtaining the data given the signal, $\Pr(s)$ is the prior probability, and the evidence $\Pr(d)$ can be considered merely as a normalisation constant. Thus, in order to obtain the (Bayesian estimator of the signal vector, we must maximise the product the product $\Pr(d|s) P(s)$ of the

likelihood function and the prior. A thorough discussion of Bayesian analysis techniques is given by Sivia (1996).

The likelihood function describes the statistics of noise contribution ϵ to the data. This function may take any form appropriate to the noise statistics. It is convenient to define the log-likelihood function $L(s) = \ln[\Pr(d|s)]$ so that the likelihood function may be written as $\Pr(d|s) = \exp[L(s)]$. As an example, if the noise on the data is Gaussian-distributed and described by the noise covariance matrix N , then the likelihood function takes the form

$$\begin{aligned} \Pr(d|s) &\propto \exp\left[-\frac{1}{2}\epsilon^T N^{-1}\epsilon\right] \\ &\propto \exp\left[-\frac{1}{2}(d - \Phi(s))^T N^{-1}(d - \Phi(s))\right], \end{aligned} \quad (2)$$

where in the second line we have used (1). In this case, the log-likelihood function is simply minus one half of the standard χ^2 misfit statistic, i.e.

$$L(s) = -\frac{1}{2}\chi^2(s).$$

The prior distribution $\Pr(s)$ codifies our knowledge of the underlying signal before acquiring the data. If we have some advance knowledge of the statistical properties of the signal then this is contained in the prior. For example, if we assume the signal to be described by a Gaussian random field with a covariance matrix C , then the prior takes the form

$$\Pr(s) \propto \exp\left(-\frac{1}{2}s^T C^{-1}s\right).$$

Indeed, if the prior is assumed to have this form and the likelihood is also Gaussian, as in (2), then the Bayesian estimator \hat{s} is obtained by maximising their product is identical to the standard Wiener filter solution. An introduction to the Wiener filter technique is given by Press et al (1994).

It is clear, however, that although the noise contribution to the data may often be Gaussian-distributed, the assumption of a Gaussian form for the prior is not valid for a general signal. If the joint probability distribution of the elements of the signal vector is known then it should be used as the prior. This is almost always impossible, however, and we instead investigate the assignment of a prior applicable to general signals that is based on information-theoretic considerations alone. Using very general notions of subset independence, coordinate invariance and system independence, it may be shown that the prior probability $\Pr(s)$ should take the form

$$\Pr(s) \propto \exp[\alpha S(s, m)], \quad (3)$$

Where the dimensional constant α depends on the scaling of the problem and may be considered as a regularising parameter, and m is a model to which the Bayesian reconstruction defaults in the absence of any data and is usually set to a small constant value. The function $S(s, m)$ is the cross-entropy of the signal and model vectors and is given by

$$S(s, m) = \sum_{n=1}^{N_s} \left[s_n - m_n - s_n \ln \left(\frac{s_n}{m_n} \right) \right], \quad (4)$$

Where N_s is the length of the (digitised) signal vector. A derivation of this result is given by Skilling (1988). By combining the entropic prior with the likelihood function, the Bayesian estimator of the signal is found by maximising with respect to s the posterior probability, which now takes the form

$$\Pr(s|d) \propto \exp[L(s)] \exp[\alpha S(s, m)] = \exp[L(s) + \alpha S(s, m)].$$

Thus, maximising this probability distribution is equivalent to maximising the function

$$F(s) = L(s) + \alpha S(s, m), \quad (5)$$

and this forms the basis of the maximum entropy method (MEM).

5

The maximum-entropy method has been applied to a wide range of signal reconstruction problems. In its standard form, however, it can be very computationally intensive. The function $F(s)$ is in general a complicated function of the components s_n of the signal vector and so a numerical maximisation of the $F(s)$ must be performed over this N_s -dimensional space. It is not unusual for N_s to be of the order $N_s \sim 10^6$, particularly in image reconstruction problems. Moreover, the standard MEM approach is only applicable to signals that are strictly non-negative, as is clear from the presence of the logarithmic term in the expression (4) for the entropy.

Nevertheless, it is possible to extend the MEM so that it can be applied to the reconstruction of signals that can take both positive and negative values. The definition of the entropy for positive/negative signals with certain special properties was first presented by Gull & Skilling (1990). The generalisation to arbitrary positive/negative signals and the derivation of the prior probability in this case is given by Hobson & Lasenby (1998). It is found that the prior has the same form as given (3) but the expression for the entropy $S(s, m)$ must be modified. The central idea is to express the general (positive/negative) signal vector s as the difference of two vectors u and v containing only strictly non-negative distributions, i.e.

$$s = u - v$$

By applying continuity constraints on the entropy functional it is possible that the expression for the entropy for the positive/negative signal s is given by

$$S(s, m_u, m_v) = \sum_{n=1}^{N_s} \left\{ \psi_n - (m_u)_n - (m_v)_n - s_n \ln \left[\frac{\psi_n + s_n}{2(m_u)_n} \right] \right\}, \quad (6)$$

Where m_u and m_v are separate models for u and v respectively, and where $\psi_n = [s_n^2 + 4(M_u)_n(M_v)_n]^{1/2}$. We cannot hope to replace the models m_u and m_v by a single positive/negative model m_s (say), since such a distribution could no longer be considered as an integration measure. Nevertheless, we can still consider the difference $m_u - m_v$ as the model for the signal s . We note that the form of the positive/negative entropy derived by Gull & Skilling (1990) requires $m_u = m_v$.

Given the entropic prior for general positive/negative signals, it is then straightforward to define the prior for complex signals simply by applying the above analysis to the real and imaginary parts separately. In this case the models M_u and M_v are also taken to be complex. The real and imaginary parts of m_u are the models for the positive portions of the real and imaginary parts of s respectively. Similarly, the real and imaginary parts of m_v are the models for the negative portions of the real and imaginary parts of the image. The total entropy of the complex signal is then obtained by evaluating the sum (6) using first the real parts and then the imaginary parts of s , m_u and m_v and adding the results.

The ability to reconstruct positive/negative and complex distributions using the MEM approach has profound consequences for greatly improving the both the speed and accuracy of the MEM technique. These improvements are based on the idea of making a change of basis in both the signal and data spaces and performing a Bayesian reconstruction in this new basis. With an appropriate choice of the new bases, it is possible to speed up significantly the calculation of the reconstruction, which can then easily be rotated back into the original basis in signal space to obtain the final reconstruction of the signal. For example, the ability to reconstruct

complex signals allows us to perform reconstructions in the Fourier basis of complex exponentials, which greatly reduces the computational complexity of de-blurring images that have been convolved with a spatially varying point-spread function (see the example below).

In order to understand how a general change of basis is performed we must first remind ourselves of some basic results in linear algebra and vector (see e.g. Riley, Hobson & Bence 1997). Suppose there exists a set of linearly independent vectors $e^{(n)}$ ($n=1, \dots, N_s$), that form a complete basis for the N_s -dimensional space of the signal vector. We may then write the signal vector as a weighted sum of these vectors. Formally, if we take $e^{(n)}$ to be the column vector with unity as the n th element and zeros elsewhere, then we may write the signal vector as

$$s = \sum_{n=1}^{N_s} s_n e^{(n)}.$$

Thus we see that in order to reconstruct the signal vector, we are in fact reconstructing its coefficients in this trivial basis. We can, however, equally well expand the signal vector in terms of any other linearly independent set of vectors $e'^{(n)}$ ($n=1, \dots, N_s$) such that

$$s = \sum_{n=1}^{N_s} s'_n e'^{(n)}.$$

25

We may perform a similar procedure in the N_d -dimensional data space, which is in general distinct from the signal space. If we consider the trivial basis vectors $f^{(n)}$ ($n=1, \dots, N_d$) in this space, with unity as the n th elements and zeros elsewhere, then the data vector is given by

$$d = \sum_{n=1}^{N_d} d_n f^{(n)}.$$

On performing change of basis in data space to some other basis $f'^{(n)}$, this becomes

$$d = \sum_{n=1}^{N_d} d'_n f'^{(n)}.$$

Since the noise vector ϵ also belongs to the data space a similar change of basis can apply to it, such that

$$\epsilon = \sum_{n=1}^{N_d} \epsilon_n f^{(n)} = \sum_{n=1}^{N_d} \epsilon'_n f'^{(n)}.$$

5

It is clear that, even if the element s_n of the signal vector in the original basis were strictly non-negative, the elements s'_n in the new basis will in general take both positive and negative values.

10 Furthermore, in the case where the new basis vectors have complex components, the coefficients s'_n may themselves be complex. Hence it is the extension of the MEM technique to the reconstruction of such quantities that allows this approach to be taken.

15 Once we have performed the changes of basis in the data and signal spaces, we denote the vectors with components s'_n by s' and we similarly define the vectors d' and ϵ' as those containing the elements d'_n and ϵ'_n respectively. In the signal space we can relate the two
20 bases $e^{(n)}$ and $e'^{(n)}$ ($n=1, \dots, N_s$) by

$$\epsilon'^{(n)} = \sum_{i=1}^{N_s} U_{in} \epsilon^{(i)}$$

where U_{in} is the i th component of $e'^{(n)}$ with respect to the original (unprimed) basis. The vector s' and s
25 are then related by

$$s = Us' \quad (7)$$

Similar results hold for the bases $f^{(n)}$ and $f'^{(n)}$ ($n = 1, \dots, N_d$) in data space, such that

$$d = Vd' \quad (8)$$

30 Where the element V_{in} is the i th component of $f'^{(n)}$ with respect to the unprimed basis. A similar expression exists relating the noise vectors ϵ' and ϵ . Substituting (7) and (8) and (1), we then obtain

$$35 \quad d' = \Phi'(s') + \epsilon', \quad (9)$$

Where ϕ' is a new function relating the signal and data vectors in the new basis.

Clearly, our aim is to choose the new bases in the data and signal spaces in order that the relationship (9) takes the simplest form. More formally, we wish to perform similarity transformations in the data and signal spaces that partition each space into numerous (quasi-) disjoint subspaces of much lower dimensionality. In such bases, we may then calculate the estimate \hat{s}' of the rotated signal vector by performing numerical maximisations in each subspace separately. Thus, we replace a single maximisation over the N_s -dimensional signal space in the original basis by numerous low-dimensionality maximisations over each subspace in the new basis. This leads to a considerable increase in the speed with which the reconstruction can be performed. Then, having calculated the Bayesian reconstruction \hat{s}' in the new basis, the required signal reconstruction \hat{s} can be obtained by rotating back to the original signal-space basis. We reiterate that, since the elements S_m in the new basis can in general take positive, negative or even complex values, it is the extension of the MEM technique to the reconstruction of such quantities that allows this approach to be taken.

As an example, let us consider the application of the above Bayesian reconstruction technique to the deconvolution of images that have been blurred by a spatially-invariant point-spread function (PSF) and that may also contain some noise contribution. For simplicity, we will assume that the de-blurred reconstruction is produced on the same pixelisation as the blurred image, although this is clearly not required by the technique in general. In this simple case, the data and signal spaces coincide.

It is well known that the convolution of an underlying image with a spatially invariant PSF is equivalent to multiplying together their Fourier

transforms and then performing an inverse Fourier transform. Therefore, in the Fourier domain, each Fourier mode can be considered independently of the others. This suggests that we should perform the Bayesian reconstruction in the Fourier basis, such that

$$s'_n = \sum_{m=1}^{N_s} \exp(-i\pi n(m-1)/N_s) s_m.$$

With a similar expression relating the components of the data vectors d' and d and the noise vectors e' and e (Since the data and signal spaces coincide). Thus, in this case, the N_s -dimensional signal (and data) space has been partitioned into N_s separate disjoint spaces (i.e. one for each Fourier mode).

Now for each value of n (or Fourier mode), we may consider the elements d'_n , s'_n and e'_n independently of those for other values of n . This leads to a substantial decrease in the CPU time required to de-blur a given image. For simplicity, at our chosen Fourier mode we denote d'_n , s'_n and e'_n by d' , s' and e' respectively. The quantity d' is given simply by the Fourier coefficient of the true underlying image, or signal s' , multiplied by the Fourier coefficient of the PSF, or response R . In addition, a noise contribution, e' , in the Fourier domain may also be present. If no instrumental errors are expected from a given apparatus, it is still possible to introduce "noise" by, for example, digitising an image in order to store it on a computer. Thus the data value is given by

$$d' = R s' + e' \quad (10)$$

Since we are performing the reconstruction in the Fourier basis, the noise on each Fourier mode will contain contributions from a wide range of scales. Therefore, provided the noise on the image is distributed in a statistically-homogeneous manner, we would expect from the central limit theorem that the noise in the Fourier domain is described reasonably well by a Gaussian

distribution. Therefore, the likelihood function is given by

$$\Pr(d'|s') \propto \exp(-|d' - Rs'|^2/\sigma^2),$$

5 where $\sigma^2 = \langle \epsilon' \epsilon'^* \rangle$ is the variance of the noise contribution at the particular Fourier mode under consideration. From (6), the entropy $S(s', m)$ of this complex "image" may be shown to be given by (where we have set $m_u = m_v = m$)

$$10 \quad S(s, m) = \Re(\psi) - 2\Re(m) - \Re(s') \ln \left[\frac{\Re(\psi + s')}{2\Re(m)} \right] + \Im(\psi) - 2\Im(m) - \Im(s') \ln \left[\frac{\Im(\psi) + \Im(s')}{2\Im(m)} \right],$$

where the \Re and \Im denote the real or imaginary part respectively of a complex number; also $\Re(\psi) = [\Re(s')^2 + 4\Re(m)^2]^{1/2}$ and a similar expression exists for $\Im(\psi)$.

15 Using the above expression for the likelihood and prior, and assuming a particular value for the regularising parameter α in (5), it is then possible numerically to maximise the posterior probability to obtain the estimator \hat{s}' of the signal vector at each Fourier mode independently. Once these estimators have been calculated for all the Fourier modes, we simply perform an inverse Fourier transform to recover the de-blurred image, i.e.

$$25 \quad \hat{s}_n = \frac{1}{N_s} \sum_{m=1}^{N_s} \exp(+i\pi n(m-1)/N_s) \hat{s}'_m.$$

The value of α used in the reconstruction algorithm may be set in three different ways. Firstly, α may be set such that the misfit statistic χ^2 , between the observed data and that predicted from the reconstruction, is equal to its expectation value, i.e. the number of data points to be fitted. This choice is usually referred to as "historic" MEM. Alternatively, it is possible to determine the appropriate value for α in a fully Bayesian manner (Skilling 1989) by simply treating it as another parameter in our hypothesis space. This is the recommended choice and is usually referred to as

"classic" MEM. Finally, the simplest option is to fix α in advance to some value. This option is unlikely to yield optimal results. It can, however, be used to obtain a quick solution if the historic or classic value for α has already been determined for a particular problem on a previous occasion.

The above technique has been applied to several different data sets in which an image has been convolved with a spatially-invariant point-spread function. For example, the "classic" de-blurring of a microscope image of a section through a pollen grain. The section has dimensions 128 x 128 and is taken from a three dimensional "stack" of 20 such images. The original stack had been blurred by a spatially invariant three dimensional PSF and the de-blurred reconstruction of the entire stack required approximately 45 seconds on an Intel Pentium 200 MHz processor. A standard MEM algorithm, which does not use the similarity transformation technique was also applied to this stack of images. This produced reconstructions of a similar quality to those obtained using the invention, but required approximately 50 minutes CPU time on the same machine.

Similar gains in speed can be obtained by expanding the signal in different bases appropriate to the given problem, so long as the correlation matrix of the resulting coefficients of the data, signal and noise vectors in the new bases is relatively sparse. For example, similar results may be obtained by reconstructing the coefficients in a wavelet expansion (Daubechie 1992) of the signal as opposed to the Fourier expansion used above. This case has the advantage that the coefficients are always real. Furthermore, the scaling/translation properties of the wavelet transform allow automatic multi-resolution reconstructions of the signal vector.

Clearly, the general method outlined above can be applied to numerous different reconstruction problems of arbitrary dimensionality. Examples include the analysis of acoustic data, radar, underwater sonar, spectroscopy, geophysical data, oil exploration and medical imaging. In addition to spatial dimensions, additional dimensions such as time, spectral behavior and polarisation are also easily accommodated.

CLAIMS

1. A method of reconstructing a signal from a given set of data, with a prediction function representing a predictable effect on the signal, and a noise function representing unpredictable noise, the method comprising the steps of:
 - altering the coordinate basis of the data and signal from an original coordinate basis in order to produce a prediction function having a reduced set of variables;
 - performing a Bayesian reconstruction capable of operation of positive, negative, and complex signal values to produce a reconstruction signal; and
 - converting the reconstruction signal back into the original coordinate basis to generate a signal.
2. A method according to claim 1, wherein the Bayesian reconstruction is performed using a Fourier basis.
3. A method according to claim 1, wherein the Bayesian reconstruction is performed using a wavelet basis.
4. A method according to any preceding claim, wherein the Bayesian reconstruction employs the maximum entropy method.
5. A method according to claim 4, employing an evaluation parameter, α , which is determined from a prior reconstruction.
6. A method according to claim 4, employing an evaluation parameter, α , which is set at a fixed value.
7. A method according to claim 4, employing an evaluation parameter, α , which is determined during the reconstruction step.

8. A method according to any of claims 1 to 7, in which the signal to be reconstructed is an image signal.
9. A method according to claim 8, wherein the image
5 signal is a medical image signal.
10. A method according to any of claims 1 to 7, in which the signal to be reconstructed is a radar signal.
- 10 11. A method according to any of claims 1 to 7, in which the signal to be reconstructed is an acoustic data signal.
12. A method according to claim 11, wherein the acoustic data signal is an underwater sonar signal.
- 15 13. A method according to claim 11, wherein the acoustic data signal is a geophysical data signal.
14. A method according to any of claims 1 to 7, in which
20 the signal to be reconstructed is a signal from spectroscopy.
15. A method according to any one of claims 1 to 7, in which the signal is a communication signal, such as a
25 time-series signal.

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 99/02715

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F17/10

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	ZHAOJING ZHOU ET AL: "A FAST MEM ALGORITHM FOR HIGH RESOLUTION DIRECTION SPECTRUM ESTIMATION" MEASUREMENT, vol. 18, no. 3, July 1996 (1996-07), pages 159-167, XP000639117 ---	
A	HOBSON M P ET AL: "The entropic prior for distributions with positive and negative values" MONTHLY NOTICES OF THE ROYAL ASTRONOMICAL SOCIETY, 11 AUG. 1998, BLACKWELL SCIENCE FOR R. ASTRON. SOC, UK, vol. 298, no. 3, pages 905-908, XP002094232 ISSN 0035-8711 --- -/--	

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

13 October 1999

Date of mailing of the international search report

20/10/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Pierfederici, A

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 99/02715

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>BURCH S F ET AL: "Image restoration by a powerful maximum entropy method"</p> <p>COMPUTER VISION, GRAPHICS, AND IMAGE PROCESSING, AUG. 1983, USA, vol. 23, no. 2, pages 113-128, XP002094233</p> <p>ISSN 0734-189X</p> <p style="text-align: center;">-----</p>	